

Edge-Based Noise Cleaning of Chrominance Information in Color Images

*James E. Adams, Jr. and John F. Hamilton, Jr.
Eastman Kodak Company
Rochester, New York*

Abstract

It is well known that noise cleaning a full-color image, as if it were a set of independent single-channel images, produces suboptimal results. Instead, it is generally recognized that transforming a color image into a luminance-chrominance space permits the use of simple and aggressive chrominance noise-cleaning operations that minimize degradation of luminance information. These chrominance noise-cleaning techniques usually consist of simple blurring operations that treat all but the lowest chrominance modulations as unwanted noise. While this is effective for noise cleaning high-frequency noise and aliasing artifacts, it can also cause color to bleed across sharp, colored edges in the image. This paper describes a noise-cleaning technique that incorporates edge detection into the chrominance blurring operation in order to prevent color bleeding at image edges. The approach is to first create a map of edge activity in the image and, during noise cleaning, to adaptively increase the noise-cleaning support region for each pixel location based on the local edge activity. By protecting edge fidelity in this manner, noise-cleaning techniques that might otherwise be too aggressive may be applied to the chrominance information. In particular, larger support regions can be used in portions of the image with lower spatial activity. Not only does this increase the effectiveness of reducing colored, high-frequency noise in an image, but it also permits the reduction of lower spatial frequency aliasing artifacts.

Introduction

One type of noise found in digital camera images appears as low-frequency, colored blobs in regions of low spatial frequency, for example, a person's face. See Fig. 1. These blobs, a type of chroma noise, produce a mottled appearance in an otherwise spatially flat region. These colored blobs are irregularly shaped and are typically 5 to 25, or more, pixels wide in a given direction.

There are numerous existing ways for reducing chroma noise in digital images. Among these are numerous patents that describe chroma noise reduction methods using optical blur filters.¹⁻³ These devices frequently address only high frequency chroma noise and are generally ineffective against low frequency chroma noise.




Figure 1. Example of chroma noise.

Another very common approach to chroma noise reduction is to use standard grayscale image noise reduction techniques on each color channel of the image, in effect, treating each color channel as a separate grayscale image.⁴ By treating a full-color image as three, unrelated grayscale images, any interactions or correlations between the color channels are ignored. As discussed below, the inherent relationships between the color planes of a digital image can be used to perform more effective chroma noise cleaning, for example, by transforming the image into a different color space that permits an easier separation of image noise from genuine scene content.

Some approaches deal specifically with digital image processing methods for reducing or removing chroma noise artifacts. One class of digital camera patents discloses improvements to the color filter array (CFA) interpolation operation to reduce or eliminate high frequency chroma noise artifacts.⁵ Another class of patents, teaches using different pixel shapes (that is, rectangles instead of

squares) and arrangements (for example, each row is offset by half a pixel width from the preceding row) with accompanying CFA interpolation operations to reduce or eliminate chroma noise artifacts.⁶ Again, these techniques address only high frequency chroma noise and are generally ineffective against low frequency chroma noise.

There is the well-known technique in the open literature of taking a digital image with chroma noise artifacts, converting the image to a luminance-chrominance space, such as CIELAB, blurring the chrominance channels and converting the image back to the original space.^{7,8} This operation is a standard technique used to combat chroma noise. One liability with this approach is that there is no discrimination during the blurring step between chroma noise artifacts and genuine chroma scene detail. Consequently, sharp colored edges in the image begin to bleed color as the blurring becomes more aggressive. See Fig. 2.




Figure 2. Example of color bleed as a result of over-aggressive chroma channel blurring: (left) before blurring, (right) after blurring

Usually, the color bleed has become unacceptable before most of the low-frequency, colored blobs are removed from the image. Also, if any subsequent image processing is performed on the image, there is the possibility of amplifying the visibility of the color bleeding. As a consequence, a small, fixed blur kernel is almost required to try to contain the problem of color bleeding. However, to address low-frequency chroma blobs, large blur kernels would be needed to achieve the desired noise cleaning.

Problem Definition

Given the existing body of work, it was determined that a chroma noise reduction method was needed that permits the use of large blur kernels while not causing color bleeding at sharp-colored edges. In order to achieve this goal, it was decided to use variable shaped (i.e., adaptive) blur kernels. These kernels would respond to edges within the image and dynamically change their shapes to avoid blurring edge boundaries. As a consequence, color bleed would be avoided while still permitting the use of large area blurring operations to eliminate low-frequency chroma noise artifacts. Of course, high-frequency chroma noise artifacts would also be eliminated. Finally, it was desired to create a method that did not require user intervention, aside from possibly setting the overall chroma noise reduction aggressiveness.

Technical Details

The resulting chroma noise-cleaning algorithm is diagrammed in Fig. 3. The first step is to convert the initial image, which is assumed to be in a standard RGB color space, to a luminance-chrominance space, such as CIELAB. Although CIELAB is the preferred working color space, other luminance-chrominance color spaces have been tried with good success. From this converted “luma-chroma” image, three-edge maps are created for each channel in the image. To create each map, four-edge detector filters are convolved with each channel and the results summed. The four filters, given in Fig. 4, are h for horizontal, v for vertical, s for slash, and b for backslash. The resulting edge map channel, $g(x)$, is computed from the image channel, $f(x)$, using Eq. 1.

$$g(x) = |h ** f(x)| + |v ** f(x)| + |s ** f(x)| + |b ** f(x)| \quad (1)$$

where x is either the luminance channel, L^* , or one of the chrominance channels, a^* or b^* , “**” is the two-dimensional convolution operation, and absolute values of the components are added together. The four-edge detector kernels are 5 x 5 truncated pyramid filters⁹ that were chosen to provide some robustness when used with noisy data. Larger kernels could be used for more noise suppression.

Once the edge maps have been created, the final composite edge map, g , is created by summing together the three channel edge maps, as given in Eq. 2.

$$g = g(L^*) + g(a^*) + g(b^*) \quad (2)$$

Because the individual components are already positive values, there is no need to use additional, absolute value operations in conjunction with this summation.




Figure 3. Chroma noise-reduction algorithm.

From the composite edge map, a threshold value is determined. Typically, for an initial estimate, one chooses a flat region of the image and calculates the standard deviation (σ) of the edge map values in this region. Three times the standard deviation (3σ) becomes the threshold. This threshold will adjust the aggressiveness of the noise cleaning, and the user could be given the ability to manually adjust this value. It was found for the

applications considered, that the threshold did not significantly change from image to image. As a result, the threshold was determined once, off-line, and, next, left as a fixed value.

The chrominance channels, a^* and b^* , are smoothed. First, the pixel to be cleaned is called the “reference pixel” and, the “reference value” is set to the corresponding edge map value for the reference pixel. Next, the algorithm moves out in each of the eight compass directions, N, NE, E, SE, S, SW, W, NW, one pixel at a time, examining the neighboring edge map values. If the difference between an edge map value and the reference value is less than the threshold, that pixel is added to the smoothing neighborhood region, and the algorithm continues. Once an edge map value is reached that differs from the reference value by more than the threshold, the growth of the smoothing neighborhood region in that compass direction is stopped. Figure 5 illustrates a typical smoothing neighborhood region around the reference pixel A, after all compass directions have been examined. In this arrangement, eight directions are shown because there are eight contiguous pixels surrounding the pixel of interest. Note that the shape and size of this neighborhood will vary from pixel to pixel. Returning to Fig. 3, once the smoothing neighborhood region is defined, the a^* and b^* channel values are averaged within this neighborhood and these averages are the noise-cleaned values for the reference pixel. Each pixel in the image is processed in this way.

$$h = \frac{1}{13} \begin{pmatrix} 1 & 1 & 0 & -1 & 1 \\ 1 & 2 & 0 & -2 & -1 \\ 1 & 2 & 0 & -2 & -1 \\ 1 & 2 & 0 & -2 & -1 \\ 1 & 1 & 0 & -1 & -1 \end{pmatrix} \quad v = \frac{1}{13} \begin{pmatrix} -1 & -1 & -1 & -1 & -1 \\ -1 & -2 & -2 & -2 & -1 \\ 0 & 0 & 0 & 0 & 0 \\ 1 & 2 & 2 & 2 & 1 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix}$$

$$s = \frac{1}{13} \begin{pmatrix} -1 & -1 & -1 & -1 & 0 \\ -1 & -2 & -2 & 0 & 1 \\ -1 & -2 & 0 & 2 & 1 \\ -1 & 0 & 2 & 2 & 1 \\ 0 & 1 & 1 & 1 & 1 \end{pmatrix} \quad b = \frac{1}{13} \begin{pmatrix} 0 & -1 & -1 & -1 & -1 \\ 1 & 0 & -2 & -2 & -1 \\ 1 & 2 & 0 & -2 & -1 \\ 1 & 2 & 2 & 0 & -1 \\ 1 & 1 & 1 & 1 & 0 \end{pmatrix}$$

Figure 4. Edge detection kernels used for creating edge maps.

The final step in Fig. 3 is to convert the noise-cleaned image back to the original color space.

As a final note, the maximum possible direction of neighborhood region expansion in any given compass direction can be restricted to prevent huge neighborhood regions from being used in large-flat regions. Such a maximum radius value might be in the range from 10 to 20 pixels.




Figure 5. Typical adaptive chroma noise-cleaning neighborhood.




Figure 6. Example of chroma noise-reduction algorithm: (left) original, (right) noise cleaned.

Discussion

Figure 6 demonstrates the effects of the algorithm. Note how integrity of the colored edges is preserved while the chroma noise is reduced in visibility.

Conclusions

A new algorithm for cleaning low-frequency chroma noise artifacts is presented. This method uses adaptively sized and shaped neighborhoods to simultaneously reduce chroma noise and preserve genuine, colored edge integrity. This algorithm presents an improvement over existing chroma noise-cleaning algorithms because it permits the use of large blur kernels while not causing color bleeding at sharp colored edges. Both high- and low-frequency chroma noise artifacts are reduced in visibility by this algorithm. Finally, this method permits the user to set the overall noise-reduction aggressiveness.

References

1. Mark M. Meyers, U.S. Patent 5,682,266 (1997).
2. Jeffery I. Hirsh, Sharlene A. Wilson, U.S. Patent 5,714,284 (1998).
3. Russell J. Palum, U.S. Patent 6,326,998 B1 (2001).
4. William K. Pratt, *Digital Image Processing*, Second Edition, John Wiley & Sons, Inc., New York, 1991, pg. 311.
5. John A. Weldy, Jennifer C. Loveridge, U.S. Patent 6,188,804 B1 (2001).
6. Michael C. Moorman, Robert H. Hibbard, Kenneth A. Parulski, U.S. Patent 5,251,019 A1 (1993).
7. <http://www.camerabits.com/techtalk.html>.
8. <http://ppgf.com/gce/jul2002/noise.html>.
9. William K. Pratt, *Digital Image Processing*, Second Edition, John Wiley & Sons, Inc., New York, 1991, pg. 504.

Biography

James E. Adams, Jr. is a Senior Principal Scientist for Eastman Kodak Company. His technical interests are in digital camera image processing, and he has 16 patents in color filter array interpolation and other digital camera image processing operations. He holds an MS in Optics from the University of Rochester.

John F. Hamilton, Jr. is a Research Fellow for Eastman Kodak Company. He is a member of Kodak's Distinguished Inventors Gallery with 25 patents in image processing algorithms for digital cameras and related applications. He holds a Ph.D. in Mathematics from Indiana University.